arXiv:1512.02394v1 [cs.LG] 8 Dec 2015

Online Gradient Descent in Function Space

Changbo Zhu, Huan Xu
National University of Singapore
Singapore
{elezhuc, mpexuh}@nus.edu.sg

October 4, 2016

Abstract

In many problems in machine learning and operations research, we need to optimize a function whose input is a random variable or a probability density function, i.e. to solve optimization problems in an infinite dimensional space. On the other hand, online learning has the advantage of dealing with streaming examples, and better model a changing environment. In this paper, we extend the celebrated online gradient descent algorithm to Hilbert spaces (function spaces), and analyze the convergence guarantee of the algorithm. Finally, we demonstrate that our algorithms can be useful in several important problems.

1 Introduction

Regret minimization is a general setting used in decision making and prediction. A merge of convex optimization and regret minimization leads to the online convex optimization problem [12]. Due to its simple setting and generality, online convex optimization has raised many attentions recently [12, 22]. Online convex optimization can be modeled as a game: at each time t, the online player chooses a point x_t from $\mathbb{K} \subset \mathbb{R}^n$. Typically, we assume that $\mathbb{K}$ is nonempty, closed and convex. After committing to this choice x_t, a convex cost function f_t is revealed and the player incurs a cost $f_t(x_t)$. Suppose this game has in total T rounds, we are interested in minimizing the regret – the gap between the actual cost and the cost of the best fixed decision in hindsight:

$$Regret(T) = \sum_{t=1}^{T} f_t(x_t) - \min_{x\in\mathbb{K}} \sum_{t=1}^{T} f_t(x).$$

Zinkevich proposed the following algorithm called online gradient descent [25] for the above problem: play $x_1 \in \mathbb{K}$ arbitrarily and at round t play $x_t := \mathbf{P}_{\mathbb{K}}(x_{t-1} - \nabla f(x_{t-1}))$, where $\mathbf{P}_{\mathbb{K}}$ stands for projection into $\mathbb{K}$. The regret of online gradient descent is shown upper bounded by $O(\sqrt{T})$. By assuming that all f_t has second derivatives bounded below by a strictly positive number, Hazan et al [13] gave an algorithm which can achieve $O(\log(T))$ regret bound. Subsequently, Hazan et al [14] provide an algorithm achieving rates inteplay between $O(\sqrt{T})$ and $O(\log(T))$ without a priori knowledge of the lower bound on the second derivatives.

In the above setup, the online player chooses a point in a *finite dimensional Euclidean space* at each step. However, in many cases, we need to make decisions over a set of functions or random variables. For example, minimizing a risk measure over a set of random variables (chapter 4 in [6]), solving an optimization problem in the space of distribution functions [11] and learning a classifier over a set of functions [16, 18, 19]. In all these cases, we need to consider optimization problems in some infinite dimensional space.

In this paper, we propose the online functional gradient algorithm which extends online gradient descent algorithm of Zinkevich [25] to a general Hilbert space (function space). We then establish the regret bound of the algorithm. Since sometimes exact projection to a convex closed set in infinite dimensional space may be hard compute, we also analyze the noisy version of the algorithm when the projection at each step is not accurate. Finally, we illustrate applications of

our algorithms in online classifier selection, risk measure minimization and distributionally robust stochastic program.

Notation: in this paper, we use $\mathcal{H}$ to denote a Hilbert space with inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ and induced norm $\|x\|_{\mathcal{H}} = \sqrt{\langle x, x \rangle}$ for any $x \in \mathcal{H}$. Also, let $\mathcal{B} = \{x \in \mathcal{H} \,|\, \|x\|_{\mathcal{H}} \leq 1\}$ denote the clsoed ball of $\mathcal{H}$ and $\mathbb{R}$ denote the set of real numbers. Let (Ω, Σ, μ) be a measure space. For $1 \leq p < \infty$, the space $\mathcal{L}^p(\Omega, \Sigma, \mu)$ is the set of all μ-measurable functions $f : \Omega \to [-\infty, +\infty]$ such that $\int |f|^p \, d\mu < \infty$. For a random variable ξ, we use small p to denote its probability density function and big P to denote its probability distribution. Further, for a sequence $\{\epsilon_t\}_{t=1}^T$, we use $\bar{\epsilon}$ to denote its average, i.e. $\bar{\epsilon} = \sum_{t=1}^T \epsilon_t / T$; and $\bar{\Upsilon}$ to denote the average of squares, i.e. $\bar{\Upsilon} = \sum_{t=1}^T \epsilon_t^2 / T$. Finally, we use $\mathbf{P}$ to denote the projection mapping and $\mathbf{E}$ to denote the expectation of a random variable.

2 Motivating Examples

In this section, we provide some examples of problems which need either gradient descent or online gradient descent *in the function space*.

2.1 Online Classifier Selection

Assuming that we can collect data points $\{(x_t, y_t)\}_{t=1}^T$ from an unknown distribution P on $\mathbb{X} \times \mathbb{Y}$, where $\mathbb{X} \subseteq \mathbb{R}^n$ is the space of data points and $\mathbb{Y} \subseteq \mathbb{R}$ is the space of labels. Then, we can learn a classifier $f : \mathbb{X} \to \mathbb{Y}$, which can best predict the label of a data point by solving the following problem

$$\min_{f \in \mathcal{C}} \frac{1}{T} \sum_{t=1}^T l(f(x_t), y_t), \tag{1}$$

where $\mathcal{C}$ is a predefined set of classifiers and l is the loss function. If f is linear, then (1) can be reduced to a standard minimization problem (like linear SVM) in $\mathbb{R}^n$. A more interesting case is studied in [18], in which the authors considered $\mathcal{C}$ as the set of all linear combinations of finitely many base-classifiers and then apply the gradient descent algorithm in a suitable function space to solve the problem. Another example is the nonlinear SVM considered in [21, 7], which is solved by letting $\mathcal{C}$ be a certain Reproducing Kernel Hilbert Space (RKHS) and using the representer theorem [15] to convert the problem into an optimization problem in $\mathbb{R}^n$.

A natural extension of (1), which we consider in this paper, is its online version. That is, the data points are supplied sequentially, and the goal of learning is to minimize the following regret,

$$Regret(T) = \frac{1}{T} \sum_{t=1}^T l(f_t(x_t), y_t) - \min_{f \in \mathcal{C}} \frac{1}{T} \sum_{t=1}^T l(f(x_t), y_t), \tag{2}$$

where $f_t \in \mathcal{C}$ only depends on $\{x_1, \cdots, x_{t-1}, f_1, \cdots, f_{t-1}\}$. Notice that in this case, it is challenging to convert Problem (2) into a minimization problem over $\mathbb{R}^n$ even with the help of representer theorem, as the data now are supplied sequentially. Hence, it is desirable to develop an online gradient descent algorithm in some proper funtion space to solve Problem (2). We remark that the case when $\mathcal{C}$ is an entire Reproducing Kernel Hilbert Space is considered in [16], which is a significantly simpler case than the general setup we considered in this paper, as there is no projection involved.

2.2 Risk Measure Minimization

Risk Measure is used to quantify and compare uncertain outcomes, which is a central concept in decision theory [1, 20]. In this subsection, let (Ω, Σ, μ) be a probability space, i.e., Ω is a set of outcomes, σ-algebra Σ is a collection of events and $\mu(\Omega) = 1$ is a probability measure. Then a real valued random variable X is a μ-measurable function $X : \Omega \to \mathbb{R}$, which represents an uncertain outcome. We shall focus on the space $\mathcal{L}^2(\Omega, \Sigma, \mu)$, which contains all the random variables X such that $\int |X|^2 \, d\mu < \infty$ and is a Hilbert space with inner product defined as $\langle X, Y \rangle = \int XY \, d\mu$. By a risk measure ρ, we mean a function $\rho : \mathcal{L}^2(\Omega, \Sigma, \mu) \to \mathbb{R}$, which assigns a real number to each

random variable (uncertain outcome). Then, a general risk measure minimization problem can be formulated as following

$$\min_{X \in \mathcal{C}} \rho(X), \tag{3}$$

where $\mathcal{C}$ is typically a subset of $\mathcal{L}^2(\Omega, \Sigma, \mu)$.
In practice, uncertain outcomes (random variables) often result from decisions (actions) in some uncertain systems [20]. Mathematically, this can be modeled by a function $f : \mathcal{S} \to \mathcal{L}^2(\Omega, \Sigma, \mu)$, where $\mathcal{S}$ stands for the set of feasible decisions, and is a subset of some vector space $\mathcal{V}$. Then, we have $\mathcal{C} = f(\mathcal{S})$ in this case. Problem (3) is generally hard even if we can convert it to an optimization problem in the Euclidean space. On the other hand, under some conditions, we can solve or estimate the true solution by directly doing gradient descent in space $\mathcal{L}^2(\Omega, \Sigma, \mu)$.

2.3 Distributionally Robust Stochastic Program

Robust Optimization (RO) is a framework in decision making under uncertainty that has attracted fast growing attention. RO addresses decision problems in which the problem parameter is not specific but known to belong to an uncertainty set [4, 9, 24, 5]. Mathematically, robust optimization problem can be formulated as following

$$\min_{x \in \mathbb{X}} \max_{\xi \in \mathcal{S}} f(x, \xi), \tag{4}$$

where $\mathbb{X} \subseteq \mathbb{R}^m$ is the feasible set of solutions, ξ is the problem parameter and $\mathcal{S} \subseteq \mathbb{R}^n$ is the unertainty set.
If the uncertainty instead is probabilistic, and is governed by a probability distribution P, which itself is uncertain and belongs to a set of distributions $\mathcal{P}$, we get the following Distributionally Robust Stochastic Program (DRSP) aka Distributionally Robust Optimization [8, 10, 23],

$$\min_{x \in \mathbb{X}} \left(\max_{P \in \mathcal{P}} \mathbf{E}_P \left[f(x, \xi) \right] \right), \tag{5}$$

where $\mathbb{X} \subseteq \mathbb{R}^n$ and $\mathcal{P}$ is the uncertainty set consists of possible distributions of the parameter. Problem (5) is usually harder than Problem (4), because the maximization part is over a subset of some infinite dimensional space. In literature, DRSP is typically solved by exploiting special structures of the uncertainty set $\mathcal{P}$ and convert the problem to an optimization problem over Euclidean space. However, this is not always possible. In this paper, we propose a novel solution approach which depends on online gradient descent algorithm in function space. We briefly describe the idea here and defer the full analysis to Section 4.3: At each step t, if a distribution $P_t \in \mathcal{P}$ is given, then we solve a x_t which approximately minimizes the function $g(x, P_t) = \mathbf{E}_{P_t} \left[f(x, \xi) \right]$. By performing online gradient desent with respect to P, we can upper bound the term $\max_P g(\bar{x}, P)$ by $\sum_{t=1}^{T} g(x_t, P_t)/T + \delta$, where $\bar{x} = (\sum_{t=1}^{T} x_t)/T$ and δ is small. Since, each $g(x_t, P_t)$ is small by the construction of x_t, we conclude that $\max_P g(\bar{x}, P)$ is also small.

3 Online Functional Gradient Descent

In this section, we present the online functional gradient descent algorithm and its variants. All proofs are postponed to the supplementary material. We also provide a brief overview of relevant background knowledge from functional analysis in the supplementary material for the completeness. Before presenting our results, we first describe the assumptions on the set, from which the online player make decisions.

Assumption 1. *$\mathcal{K} \subseteq \mathcal{H}$ is nonempty, closed, convex and $\mathcal{K} \subseteq R\mathcal{B}$ for some $R \in \mathbb{R}$.*

Also, we make some assumptions on the cost functions received by the player.

Assumption 2. *$f : \mathcal{K} \to \mathbb{R}$ is convex, Gâteaux differentiable over $\mathcal{K}$ and all the Gâteaux gradients $\{\nabla f(x) \,|x \in \mathcal{K}\}$ have finite norm, i.e. $\|\nabla f(x)\|_{\mathcal{H}} < +\infty$ for all $x \in \mathcal{K}$.*

Lemma 1. *Let $\mathcal{C} \subseteq \mathcal{H}$ be a nonempty closed convex subset, then for all $x \in \mathcal{H}$ and $\hat{x} \in \mathcal{C}$, we have*

$$\|\mathbf{P}_{\mathcal{C}}(x) - \hat{x}\|_{\mathcal{H}} \leq \|x - \hat{x}\|_{\mathcal{H}}.$$

Proof. By Theorem 3.14 in [2], we have

$$\langle \hat{x} - \mathbf{P}_{\mathcal{C}}(x), x - \mathbf{P}_{\mathcal{C}}(x) \rangle_{\mathcal{H}} \leq 0,$$

which further implies that

$$\|\mathbf{P}_{\mathcal{C}}(x) - \hat{x}\|_{\mathcal{H}}^2 + \|x - \mathbf{P}_{\mathcal{C}}(x)\|_{\mathcal{H}}^2 \leq \|x - \hat{x}\|_{\mathcal{H}}^2.$$

□

Throughout this section, let $x^* \in \mathcal{K}$ be the optimal solution of $\sum_{t=1}^T f_t : \mathcal{K} \to \mathbb{R}$ or $f : \mathcal{K} \to \mathbb{R}$ over $\mathcal{K}$, then the online functional gradient descent algorithm proceeds as follows: pick $x_1 \in \mathcal{K}$ arbitrarily and for $t = 2, \cdots, T$, choose x_t as $x_t = \mathbf{P}_{\mathcal{K}}\left(x_{t-1} - \eta \nabla f_{t-1}\left(x_{t-1}\right)\right)$. Then, we have the following theorem upper bounds the regret.

Theorem 1. *[Online Functional Gradient Descent] Suppose $\mathcal{K}$ satisfies Assumption 1 and let $f_1, f_2, \cdots, f_T : \mathcal{K} \to \mathbb{R}$ be an arbitrary sequence of functions that satisfy Assumption 2. Pick $x_1 \in \mathcal{K}$ arbitrarily and let $x_2, \cdots, x_T$ be defined by $x_{t+1} = \mathbf{P}_{\mathcal{K}}\left(x_t - \eta \nabla f_t\left(x_t\right)\right)$. Let $G = \max_t \|\nabla f_t(x_t)\|_{\mathcal{H}}$ and select $\eta = R/G\sqrt{T}$, we have*

$$\sum_{t=1}^T f_t(x_t) - \sum_{t=1}^T f_t(x^*) \leq RG\sqrt{T}.$$

Proof. Since each f_t is convex, Gâteaux differentiable over $\mathcal{K}$, by Theorem 7.3.6 in [17], there exists a x^* in $\mathcal{K}$ minimizing $\sum_{t=1}^T f_t(x)$. Since f_t is convex, by Proposition 17.10 in [2], we have

$$f_t(x_t) - f_t(x^*) \leq \langle \nabla f_t(x_t), x_t - x^* \rangle_{\mathcal{H}}. \tag{6}$$

Since $\mathcal{K}$ is nonempty, convex and closed, by Lemma 1, we have for all $x \in \mathcal{H}$, $\|\mathbf{P}_{\mathcal{K}}(x) - x^*\|_{\mathcal{H}} \leq \|x - x^*\|_{\mathcal{H}}$. So,

$$\begin{aligned}
\|x_{t+1} - x^*\|_{\mathcal{H}}^2 &= \|\mathbf{P}_{\mathcal{K}}\left(x_t - \eta \nabla f_t(x_t)\right) - x^*\|_{\mathcal{H}}^2 \\
&\leq \|x_t - \eta \nabla f_t(x_t) - x^*\|_{\mathcal{H}}^2 \\
&= \|x_t - x^*\|_{\mathcal{H}}^2 + \eta^2 \|\nabla f_t(x_t)\|_{\mathcal{H}}^2 - 2\eta \langle \nabla f_t(x_t), x_t - x^* \rangle_{\mathcal{H}} \\
&\leq \|x_t - x^*\|_{\mathcal{H}}^2 + \eta^2 G^2 - 2\eta \langle \nabla f_t(x_t), x_t - x^* \rangle_{\mathcal{H}}.
\end{aligned}$$

After rearranging terms, we have

$$\langle \nabla f_t(x_t), x_t - x^* \rangle_{\mathcal{H}} \leq \frac{\|x_t - x^*\|_{\mathcal{H}}^2 - \|x_{t+1} - x^*\|_{\mathcal{H}}^2 + \eta^2 G^2}{2\eta}.$$

Combining with equation (6) and summing over t, we have

$$\begin{aligned}
\sum_{t=1}^T f_t(x_t) - \sum_{t=1}^T f_t(x^*) &\leq \frac{R^2}{2\eta} + T\frac{\eta G^2}{2} \\
&\leq RG\sqrt{T}.
\end{aligned}$$

□

If we only consider a single function, i.e. all the f_t are the same, then the online functional gradient descent reduces to functional gradient descent, with its performance characterized by the following corollary. The idea of gradient descent and variants in function space has appeared before [18, 19, 16], but all these works consider the unconstrained case only, i.e., $\mathcal{K} = \mathcal{H}$, and hence is a significantly easier case as no projection is involved.

Corollary 1 (Functional Gradient Descent)**.** *Suppose $\mathcal{K}$ satisfies Assumption 1 and $f : \mathcal{K} \to \mathbb{R}$ satisfies Assumption 2. Pick $x_1 \in \mathcal{K}$ arbitrarily and let $x_2, \cdots, x_T$ be defined by $x_{t+1} = \mathbf{P}_{\mathcal{K}}\left(x_t - \eta \nabla f\left(x_t\right)\right)$. Let $G = \max_t \|\nabla f(x_t)\|_{\mathcal{H}}$ and select $\eta = R/G\sqrt{T}$, we have*

$$f\left(\frac{1}{T}\sum_{t=1}^{T} x_t\right) - f(x^*) \leq \frac{RG}{\sqrt{T}}.$$

Proof. Follows from the convexity of f and Theorem 1. □

In some cases, the projection $\mathbf{P}_{\mathcal{K}}$ in general Hilbert space is not easy to calculate exactly. Hence we develop the following results with respect to noisy projection, which shows the online gradient algorithm achieves comparable guarantees if only approximated projection is used at each step.

Theorem 2 (Online Functional Gradient Descent with Noisy Projection)**.** *Suppose $\mathcal{K}$ satisfies Assumption 1 and $f_1, f_2, \cdots, f_T : \mathcal{K} \to \mathbb{R}$ be an arbitrary sequence of functions that satisfy Assumption 2. Pick $x_1 \in \mathcal{K}$ arbitrarily and let $x_2, \cdots, x_T$ be calculated such that $\|x_{t+1} - \mathbf{P}_{\mathcal{K}}\left(x_t - \eta \nabla f_t\left(x_t\right)\right)\|_{\mathcal{H}} \leq \epsilon_t$. Let $G = \max_t \|\nabla f_t(x_t)\|_{\mathcal{H}}$ and select $\eta = \sqrt{R^2/T + 4R\bar{\epsilon} + \bar{\Upsilon}}/G$, we have*

$$\sum_{t=1}^{T} f_t(x_t) - \sum_{t=1}^{T} f_t(x^*) \leq RG\sqrt{T} + \left(2\sqrt{R\bar{\epsilon}} + \sqrt{\bar{\Upsilon}}\right) GT. \tag{7}$$

Proof. Since each f_t is convex, Gâteaux differentiable over $\mathcal{K}$, by Theorem 7.3.6 in [17], there exists a x^* in $\mathcal{K}$ minimizing $\sum_{t=1}^{T} f_t(x)$. Since f_t is convex, by Proposition 17.10 in [2], we have

$$f_t(x_t) - f_t(x^*) \leq \left\langle \nabla f_t(x_t), x_t - x^* \right\rangle_{\mathcal{H}}. \tag{8}$$

Since $\mathcal{K}$ is nonempty, convex and closed, by Lemma 1, we have for all $x \in \mathcal{H}$, $\|\mathbf{P}_{\mathcal{K}}(x) - x^*\|_{\mathcal{H}} \leq \|x - x^*\|_{\mathcal{H}}$. So,

$$\begin{aligned}
\|x_{t+1} - x^*\|_{\mathcal{H}}^2 &\leq \left(\|x_{t+1} - \mathbf{P}_{\mathcal{K}}(x_t - \eta \nabla f_t(x_t))\|_{\mathcal{H}} + \|\mathbf{P}_{\mathcal{K}}(x_t - \eta \nabla f_t(x_t)) - x^*\|_{\mathcal{H}}\right)^2 \\
&\leq \|x_t - \eta \nabla f_t(x_t) - x^*\|_{\mathcal{H}}^2 + 4R\epsilon_t + \epsilon_t^2 \\
&= \|x_t - x^*\|_{\mathcal{H}}^2 + \eta^2 \|\nabla f_t(x_t)\|_{\mathcal{H}}^2 - 2\eta \left\langle \nabla f_t(x_t), x_t - x^* \right\rangle_{\mathcal{H}} + 4R\epsilon_t + \epsilon_t^2 \\
&\leq \|x_t - x^*\|_{\mathcal{H}}^2 + \eta^2 G^2 - 2\eta \left\langle \nabla f_t(x_t), x_t - x^* \right\rangle_{\mathcal{H}} + 4R\epsilon_t + \epsilon_t^2.
\end{aligned}$$

After rearranging terms, we have

$$\left\langle \nabla f_t(x_t), x_t - x^* \right\rangle_{\mathcal{H}} \leq \frac{\|x_t - x^*\|_{\mathcal{H}}^2 - \|x_{t+1} - x^*\|_{\mathcal{H}}^2 + \eta^2 G^2 + 4R\epsilon_t + \epsilon_t^2}{2\eta}.$$

Combining with equation (8) and summing over t, we have

$$\begin{aligned}
\sum_{t=1}^{T} f_t(x_t) - \sum_{t=1}^{T} f_t(x^*) &\leq \frac{R^2 + \sum_{t=1}^{T}(4R\epsilon_t + \epsilon_t^2)}{2\eta} + T\frac{\eta G^2}{2} \\
&= G\sqrt{R^2 T + \sum_{t=1}^{T} \left(4R\epsilon_t + \epsilon_t^2\right) T} \\
&\leq RG\sqrt{T} + \left(2\sqrt{R\bar{\epsilon}} + \sqrt{\bar{\Upsilon}}\right) GT.
\end{aligned}$$

□

Similarly, when all the functions $\{f_t\}$ are the same, we have the following result:

Corollary 2. *[Functional Gradient Descent with Noisy Projection] Suppose $\mathcal{K}$ satisfies Assumption 1 and $f : \mathcal{K} \to \mathbb{R}$ satisfies Assumption 2. Pick $x_1 \in \mathcal{K}$ arbitrarily and let $x_2, \cdots, x_T$ be calculated such that $\|x_{t+1} - \mathbf{P}_{\mathcal{K}}\left(x_t - \eta \nabla f\left(x_t\right)\right)\|_{\mathcal{H}} \leq \epsilon_t$. Let $G = \max_t \|\nabla f(x_t)\|_{\mathcal{H}}$ and select $\eta = \sqrt{R^2/T + 4R\bar{\epsilon} + \bar{\Upsilon}}/G$, we have*

$$f\left(\frac{1}{T}\sum_{t=1}^{T} x_t\right) - f(x^*) \leq \frac{RG}{\sqrt{T}} + \left(2\sqrt{R\bar{\epsilon}} + \sqrt{\bar{\Upsilon}}\right) G.$$

Proof. Follows from the convexity of f and Theorem 2. □

We remark that the Gâteaux differentiability used in previous results is to guarantee that the optimal solution x^* exists in $\mathcal{C}$. We can relax this assumption and use subgradient instead, which leads to the following results:

Assumption 3. *$f : \mathcal{K} \to \mathbb{R}$ is convex, subdifferentiable function over $\mathcal{K}$ and all the subgradients $\{u_x \,|x \in \mathcal{K}$ and $u_x \in \partial f(x)\}$ have finite norm, i.e. $\|u_x\|_{\mathcal{H}} < +\infty$ for all $x \in \mathcal{K}$ and $u_x \in \partial f(x)$.*

Corollary 3. *Suppose $\mathcal{K}$ satisfies Assumption 1 and $f_1, f_2, \cdots, f_T : \mathcal{K} \to \mathbb{R}$ be an arbitrary sequence of functions that satisfy Assumption 3. Pick $x_1 \in \mathcal{K}$ arbitrarily and let $x_2, \cdots, x_T$ be calculated such that $\|x_{t+1} - \mathbf{P}_{\mathcal{K}}(x_t - \eta u_t)\|_{\mathcal{H}} \leq \epsilon_t$, where $u_t \in \partial f_t(x_t)$. Let $G = \max_t \|u_t\|_{\mathcal{H}}$ and select $\eta = \sqrt{R^2/T + 4R\bar{\epsilon} + \bar{\Upsilon}}/G$, we have*

$$\sum_{t=1}^{T} f_t(x_t) - \sum_{t=1}^{T} f_t(x) \leq RG\sqrt{T} + \left(2\sqrt{R\bar{\epsilon}} + \sqrt{\bar{\Upsilon}}\right) GT, \quad \textit{for all } x \in \mathcal{K}.$$

Proof. By the definition of subgradient, it holds that

$$f_t(x_t) - f_t(x) \leq \langle u_t, x_t - x\rangle_{\mathcal{H}} \quad \text{for all } x \in \mathcal{H} \text{ and } u_t \in \partial f_t(x_t). \tag{9}$$

Then, replace equation (8) with (9) and change $\nabla f_t(x_t)$ to u_t in the rest proof of Theorem 2, we can prove the result. □

Corollary 4. *Suppose $\mathcal{K}$ satisfies Assumption 1 and $f : \mathcal{K} \to \mathbb{R}$ satisfies Assumption 3. Pick $x_1 \in \mathcal{K}$ arbitrarily and let $x_2, \cdots, x_T$ be calculated such that $\|x_{t+1} - \mathbf{P}_{\mathcal{K}}(x_t - \eta u_t)\|_{\mathcal{H}} \leq \epsilon_t$, where $u_t \in \partial f(x_t)$. Let $G = \max_t \|u_t\|_{\mathcal{H}}$ and select $\eta = \sqrt{R^2/T + 4R\bar{\epsilon} + \bar{\Upsilon}}/G$, we have*

$$f\left(\frac{1}{T}\sum_{t=1}^{T} x_t\right) - f(x) \leq \frac{RG}{\sqrt{T}} + \left(2\sqrt{R\bar{\epsilon}} + \sqrt{\bar{\Upsilon}}\right) G, \quad \textit{for all } x \in \mathcal{K}.$$

Proof. Follows from the convexity of f and Corollary 2. □

3.1 Calculate the Projection

Here, we introduce some exact formulas to calculate projections onto some common sets and refer Chapter 28 in [2] for a comprehensive study of the projection operator in Hilbert space. The following two examples show that it is easy to project a point onto a closed ball or a hyperplane.

Example 1. *Let $\mathcal{B} \subseteq \mathcal{H}$ be the closed ball with radius 1, then*

$$\forall x \in \mathcal{H} \;\; \mathbf{P}_{\mathcal{B}}(x) = \frac{1}{\max\{\|x\|_{\mathcal{H}}, 1\}} x.$$

Example 2. *Let $u \in \mathcal{H}$ be a nonzero vector, let $\eta \in \mathbb{R}$ and set $\mathcal{C} = \{x \in \mathcal{H} | \langle x, u\rangle_{\mathcal{H}} = \eta\}$, then we have*

$$\forall x \in \mathcal{H} \;\; \mathbf{P}_{\mathcal{C}}(x) = \;\; x + \tfrac{\eta - \langle x, u\rangle_{\mathcal{H}}}{\|u\|_{\mathcal{H}}^2} u.$$

Sometimes, a function $f \in \mathcal{L}^2(\Omega, \Sigma, \mu)$ is required to have nonnegative values (e.g., density functions), in which case we have the following formula:

Example 3. *Set $\mathcal{C} = \left\{p \in \mathcal{L}^2(\Omega, \Sigma, \mu) \,|\, p \geq 0\right\}$, then for any $q \in \mathcal{L}^2(\Omega, \Sigma, \mu)$, the projection $\mathbf{P}_{\mathcal{C}}(q)$ is given by*

$$\mathbf{P}_{\mathcal{C}}(q) = [q]_+, \;\; \textit{where for all } x \in \Omega, \; [q]_+(x) = \begin{cases} q(x) & \textit{if } q(x) \geq 0, \\ 0 & \textit{if } q(x) < 0. \end{cases}$$

The following algorithm explains how to project a point onto the intersection of multiple sets, if calculating the projection onto each set is easy.

Theorem 3 (Dykstra's Algorithm)**.** *Let m be a strictly positive integer, set $I = \{1, \cdots, m\}$, let $(\mathcal{C}_i)_{i\in I}$ be a family of closed convex subsets of $\mathcal{H}$ such that $\mathcal{C} = \cap_{i\in I}\mathcal{C}_i \neq \emptyset$, and let $x_0 \in \mathcal{H}$. Set*

$$i : \mathbb{N} \to I \text{ as } i(n) = 1 + rem(n-1, m),$$

where $rem(\cdot, m)$ is the remainder function of the division by m. For every strictly positive integer n, set $\mathbf{P}_n = \mathbf{P}_{\mathcal{C}_n}$, where $\mathcal{C}_n = \mathcal{C}_{i(n)}$ if $n > m$. Moreover, set $q_{-(m-1)} = \cdots = q_{-1} = q_0 = 0$ and

$$(\forall n \in \mathbb{N} \diagdown \{0\}) \quad \begin{cases} x_n = \mathbf{P}_n(x_{n-1} + q_{n-m}), \\ q_n = x_{n-1} + q_{n-m} - x_n. \end{cases}$$

Then $x_n \to \mathbf{P}_{\mathcal{C}}(x_0)$.

4 Applications

In this section, we discuss some concrete examples to illustrate how to apply the developed framework. In particular, how to compute the corresponding derivatives and the projections. For each particular application, a suitable Hilbert space is chosen.

4.1 Online Classifier Selection

The online classifier selection problem is described above in Section 2.1. Here, let $\mathcal{H}$ be a Reproducing Kernel Hilbert Space of real valued functions defined on $\mathbb{X} \subseteq \mathbb{R}^n$ and associated with a reproducing kernel $k : \mathbb{X} \times \mathbb{X} \to \mathbb{R}$. We consider the problem of minimizing the following regret:

$$Regret(T) = \frac{1}{T}\sum_{t=1}^{T} l_t(f_t) - \min_{f\in\mathcal{C}} \frac{1}{T}\sum_{t=1}^{T} l_t(f), \tag{10}$$

where $\{l_t\}$ and $\mathcal{C}$ are defined by

$$\begin{cases} l_t(f) = (f(x_t) - y_t)^2, \\ \mathcal{C} = \left(\bigcap_{i=1}^{m} \hat{\mathcal{C}}_i\right) \cap R\mathcal{B}, \text{ where } \mathcal{B} \text{ is the closed ball of } \mathcal{H}, \\ \text{Given } g_i \in \mathcal{H}, a_i \in \mathbb{R}, \text{ we have } \hat{\mathcal{C}}_i = \{ f \in \mathcal{H} | \langle f, g_i\rangle_{\mathcal{H}} = a_i \}. \end{cases}$$

Each $\hat{C}_i$ corresponds to a linear constraint imposed on f. For instance, suppose we want to guarantee that $f(x_1) = 1$ (e.g., $(x_1, 1)$ is a sample of high-confidence), we can add the linear constraint $\langle k(\cdot, x_1), f\rangle_{\mathcal{H}} = 1$. For the above problem, we can apply the online functional gradient descent algorithm, in which projection and gradient are calculated as follows:
Gradient: we first calculate the gradient of l_t for $t = 1, \cdots, T$

$$\begin{aligned} \text{for any } h \in \mathcal{H},\ \nabla l_t(f)(h) =& \lim_{\alpha\downarrow 0} \frac{l_t(f+\alpha h) - l_t(f)}{\alpha} \\ =& \lim_{\alpha\downarrow 0} \frac{(f(x_t) + \alpha h(x_t) - y_t)^2 - (f(x_t) - y_t)^2}{\alpha} \\ =& 2(f(x_t) - y_t)\, h(x_t), \\ =& \langle 2(f(x_t) - y_t)\, k(\cdot, x_t), h\rangle_{\mathcal{H}} \text{ (reproducing property).} \end{aligned}$$

By Remark 2.44 in [2], we have $\nabla l_t(f) = 2(f(x_t) - y_t)\, k(\cdot, x_t)$.
Projection: for any $g \in \mathcal{H}$, we calculate the projection $\mathbf{P}_{\mathcal{C}}(g)$ onto $\mathcal{C}$. Firstly, we can use Example 2 to calculate the projection $\mathbf{P}_{\hat{\mathcal{C}}_i}(g)$ onto each $\hat{\mathcal{C}}_i$ and Example 1 to calculate the projection $\mathbf{P}_{R\mathcal{B}}(g)$ onto $R\mathcal{B}$. Then, we can calculate the projection onto the intersection of $\hat{\mathcal{C}}_i$ and $R\mathcal{B}$, i.e. onto $\mathcal{C}$, using Theorem 3.

4.2 Risk Measure Minimization

The risk measure minimization problem is introduced in Section 2.2. Here, we consider the following mean variance risk measure minimization problem as a concrete example:

$$\min_{X\in\mathcal{C}} \mathbf{E}(X) + c\|X - \mathbf{E}(X)\|^2_{\mathcal{L}^2(\Omega,\Sigma,\mu)},$$

where $c \geq 0$ is a given constant and $\mathcal{C}$ is defined as

$$\begin{cases} \mathcal{C} = \left(\bigcap_{i=1}^m \hat{\mathcal{C}}_i\right) \cap R\mathcal{B}, \text{ where } \mathcal{B} \text{ is the closed ball of } \mathcal{L}^2(\Omega,\Sigma,\mu), \\ \text{Given } Y_i \in \mathcal{L}^2(\Omega,\Sigma,\mu), a_i \in \mathbb{R}, \text{ we have } \hat{\mathcal{C}}_i = \left\{ X \in \mathcal{L}^2(\Omega,\Sigma,\mu) \middle| \langle X, Y_i\rangle_{\mathcal{L}^2(\Omega,\Sigma,\mu)} = a_i \right\}. \end{cases}$$

Notice that each $\hat{C}_i$ is a linear constraint which stands for $\mathbf{E}(XY_i) = a_i$. In particular, if $Y_i(x) \equiv 1$ for all $x \in \Omega$, then $\hat{C}_i$ is the set of random variables whose mean is a_i. Set $\rho(X) = \mathbf{E}(X) + c\|X - \mathbf{E}(X)\|^2_{\mathcal{L}^2(\Omega,\Sigma,\mu)}$, it is proved in [6] (Chapter4 page128) that ρ is convex. Also, the derivative and projection can be calculated as following:
Gradient: we first calculate the gradient of ρ at X, for any $Y \in \mathcal{L}^2(\Omega,\Sigma,\mu)$,

$$\begin{aligned} \nabla\rho(X)(Y) &= \lim_{\alpha\downarrow 0} \frac{\rho(X+\alpha Y) - \rho(X)}{\alpha} \\ &= \lim_{\alpha\downarrow 0} \frac{\int \alpha Y\, d\mu + c\|X + \alpha Y - \int X + \alpha Y\, d\mu\|^2_{\mathcal{L}^2(\Omega,\Sigma,\mu)} - c\|X - \int X\, d\mu\|^2_{\mathcal{L}^2(\Omega,\Sigma,\mu)}}{\alpha} \\ &= \int Y\, d\mu + 2c\left\langle X - \int X\, d\mu, Y - \int Y\, d\mu \right\rangle_{\mathcal{L}^2(\Omega,\Sigma,\mu)} \\ &= \langle 1 + 2cX - 2c\mathbf{E}(X), Y\rangle_{\mathcal{L}^2(\Omega,\Sigma,\mu)}. \end{aligned}$$

So, we have $\nabla\rho(X) = 1 + 2cX - 2c\mathbf{E}(X)$.
Projection: for any $Y \in \mathcal{L}^2(\Omega,\Sigma,\mu)$, the projection onto each $\hat{\mathcal{C}}_i$ and $R\mathcal{B}$ can be calculated seperately by Example 2 and 1. Then, we can use Theorem 3 to calculate the projection onto $\mathcal{C}$.

4.3 Distributionally Robust Stochastic Program

Applying online functional gradient descent to solve Distributionally Robust Stochastic Program is more involved than the previous examples. Here, we consider a slightly different version of Problem (5), i.e., we focus on probability density functions instead of probability distributions. In particular, we consider the space $\mathcal{L}^2(\mathbb{R}^n) = \mathcal{L}^2(\mathbb{R}^n, \Sigma, \mu)$, where Σ and μ are the σ-algebra of Lebesgue measurable sets and the Lebesgue measure on $\mathbb{R}^n$ respectively. Then, $\mathcal{L}^2(\mathbb{R}^n)$ is a Hilbert space with inner product defined as $\langle f, g\rangle_{\mathcal{L}^2} = \int fg\, d\mu$. Correspondingly, the induced norm on $\mathcal{L}^2(\mathbb{R}^n)$ is $\|\cdot\|_{\mathcal{L}^2}$ and we use $\mathcal{B}_{\mathcal{L}^2}$ to denote the closed ball of $\mathcal{L}^2(\mathbb{R}^n)$. Assuming that ξ is a random variable with probability density function $p \in \mathcal{P}$, then the Distributionally Robust Stochastic Program can be written as

$$\min_{x\in\mathbb{X}} \left(\max_{p\in\mathcal{P}} \mathbf{E}_p\left[f(x,\xi)\right] \right), \tag{11}$$

where $\mathbb{X} \subseteq \mathbb{R}^n$ and $\mathcal{P} \subseteq \mathcal{L}^2(\mathbb{R}^n)$ is a set of probability density functions. The main idea to solve this problem is inspired by [3], in which they use online learning methods to solve (arguably easier) Robust Optimization.
Optimization via Binary Search: We first convert the optimization problem into a decision problem. Suppose we know the feasible range of the objective value of $\max_{p\in\mathcal{P}} \mathbf{E}_p\left[f(x,\xi)\right]$, Problem (11) can be solved by a binary search procedure in the following way: let b be our current guess of the optimal value and shift f downwards by b, i.e., $h(x,\xi) = f(x,\xi) - b$. Then, we solve the following decision problem, i.e., a YES or NO problem,

$$\exists? x \in \mathbb{R}^n \text{ such that } \mathbf{E}_p\left[h(x,\xi)\right] \leq 0 \ \ \forall p \in \mathcal{P}. \tag{12}$$

An answer YES means the true optimal value is smaller than b and NO means larger. Correspondingly, if the answer is YES (NO), we should make our new guess smaller (bigger). In the rest of the section we will focus on solving Problem (12). We say that x is a δ-approximate solution to Problem (12) if $\mathbf{E}_p[h(x,\xi)] \leq \delta$ for all $p \in \mathcal{P}$. In the following, we will use online functional gradient descent with noisy projection to get a 2δ-approximate solution to Problem (12).

Oracle: notice that if we fix $p \in \mathcal{P}$, then $\mathbf{E}_p[h(x,\xi)]$ is a function mapping from $\mathbb{R}^n$ to $\mathbb{R}$. We assume this finite dimensional function is easy to optimize. In particular, we assume that there exists an Oracle $\mathcal{O}_\delta$ such that given any $p \in \mathcal{P}$, it either returns a $x \in \mathbb{X}$ such that

$$\mathbf{E}_p[h(x,\xi)] \leq \delta,$$

or return "infeasible" if there does not exist a vector $x \in \mathbb{X}$ such that

$$\mathbf{E}_p[h(x,\xi)] \leq 0.$$

4.3.1 Functional Dual Gradient Descent

Setting $g(x,p) = \mathbf{E}_p[h(x,\xi)]$, we assume that g is convex in x (which is true when $h(\cdot,\xi)$ is convex), $\mathcal{P} \subseteq R\mathcal{B}_{\mathcal{L}^2}$ is convex and $\max_{x,p} \|\nabla_p g(x,p)\|_{\mathcal{L}^2} \leq G$. Then, we propose the following algorithm to solve Problem (12).

Algorithm 1: Functional Dual Gradient Descent with Noisy Projection

input : δ, G
Initialize $p_0 \in \mathcal{P}$ arbitrarily
Choose T, $\{\epsilon_t\}_{t=1}^T$ such that $\frac{RG}{\sqrt{T}} + \left(2\sqrt{R\bar{\epsilon}} + \sqrt{\bar{\Upsilon}}\right) G \leq \delta$ and set $\eta = \sqrt{R^2/T + 4R\bar{\epsilon} + \bar{\Upsilon}}/G$
for $t = 1, 2, \cdots, T$ **do**
  Calculate p_t such that $p_t \in \mathcal{P}$ and $\|p_t - \mathbf{P}_{\mathcal{P}}(p_{t-1} + \eta \nabla_p g(x_{t-1}, p_{t-1}))\|_{\mathcal{L}^2} \leq \epsilon_t$
  Set $x_t = \mathcal{O}_\delta(p_t)$
  if *the Oracle declared infeasibility* **then**
    return *"NO"*
return $\bar{x} = \frac{1}{T}\sum_{t=1}^T x_t$

Theorem 4. *Algorithm 1 either returns an 2δ-approximate solution or concludes that the answer is "No" for Problem* (12)*.*

Proof. First, if the algorithm returns "NO", then by the definition of the Oracle, for some t, there does not exists $x \in \mathbb{R}^n$ such that $\mathbf{E}_{p_t}[h(x,\xi)] \leq 0$, which means the answer to Problem (12) is "NO". Second, otherwise, then a solution is returned, and the premise of the oracle implies that

$$\frac{1}{T}\sum_{t=1}^T g(x_t, p_t) \leq \delta.$$

From the regret guarantee of the online functional gradient descent algorithm we have

$$\max_{p\in\mathcal{P}} \frac{1}{T}\sum_{t=1}^T g(x_t,p) - \frac{1}{T}\sum_{t=1}^T g(x_t,p_t) \leq \frac{RG}{\sqrt{T}} + \left(2\sqrt{R\bar{\epsilon}} + \sqrt{\bar{\Upsilon}}\right) G \leq \delta.$$

We conclude that

$$\delta \geq \frac{1}{T}\sum_{t=1}^T g(x_t,p_t) \geq \max_{p\in\mathcal{P}} \frac{1}{T}\sum_{t=1}^T g(x_t,p) - \delta \geq \max_{p\in\mathcal{P}} g(\bar{x},p) - \delta.$$

Hence, we have $g(\bar{x},p) \leq 2\delta$ for all $p \in \mathcal{P}$. □

4.3.2 A Concrete Example

In this subsection, we consider the following simple example to illustrate how to solve DRSP via the framework outlined above:

$$\min_{x\in\mathbb{B}} \max_{p\in\mathcal{P}} \mathbf{E}_p \left[(x^T\xi)^2\right], \tag{13}$$

where $\mathbb{B} \subseteq \mathbb{R}^2$ is the closed ball of $\mathbb{R}^2$, the random variable ξ takes value from $\mathbb{C} = \left\{(x,y) \,|\, x^2+y^2 \le 1, x \ge 0, y \ge 0\right\}$ and the uncertainty set $\mathcal{P}$ is defined as

$$\mathcal{P} = \left\{p \in \mathcal{L}^2(\mathbb{R}^2) \,\middle|\, \mathbf{E}_p[\xi] = b \in \mathbb{R}^2, p \ge 0, \textstyle\int_{\mathbb{C}} p\, d\mu = 1, \int_{\mathbb{C}} |p|^2\, d\mu \le 2 \right\}.$$

We define the characteristic function $\chi_{\mathbb{C}}$ by $\chi_{\mathbb{C}}(x) = 1$ if $x \in \mathbb{C}$ and $\chi_{\mathbb{C}}(x) = 0$ otherwise. Further, set $\xi = (\xi_1, \xi_2)^T$ and $t = (t_1, t_2)^T$ with $t_i(x) = x_i$ for $x \in \mathbb{R}^2, i = 1, 2$. Then, we can write $\mathcal{P} = \mathcal{P}_1 \cap \mathcal{P}_2 \cap \mathcal{P}_3 \cap \mathcal{P}_4 \cap \mathcal{P}_5$, where

$$\begin{cases} \mathcal{P}_i = \left\{p \in \mathcal{L}^2(\mathbb{R}^2) \,|\, \mathbf{E}_p(\xi_i) = \langle t_i\chi_{\mathbb{C}}, p\rangle_{\mathcal{L}^2} = b_i\right\} \text{ for } i = 1, 2; \\ \mathcal{P}_3 = \left\{p \in \mathcal{L}^2(\mathbb{R}^2) \,|\, p \ge 0\right\}, \mathcal{P}_4 = \left\{p \in \mathcal{L}^2(\mathbb{R}^2) \,|\, \langle\chi_{\mathbb{C}}, p\rangle_{\mathcal{L}^2} = 1\right\}, \mathcal{P}_5 = 2\mathcal{B}_{\mathcal{L}^2}. \end{cases}$$

We can solve Problem (13) using Algorithm 1, in which the gradient and projection can be calculated as following:

Gradient: fix any $x \in \mathbb{R}^2$ and set $g(p) = \mathbf{E}_p\left[(x^T\xi)^2\right]$, we calculate $\nabla g(p)$ with respect to p,

$$\begin{aligned} \text{for any } q \in \mathcal{L}^2(\mathbb{R}^2),\ \nabla g(p)(q) &= \lim_{\alpha\downarrow 0} \frac{g(p+\alpha q) - g(p)}{\alpha} \\ &= \lim_{\alpha\downarrow 0} \frac{\int \chi_{\mathbb{C}}(x^Tt)^2(p+\alpha q)\, d\mu - \int \chi_{\mathbb{C}}(x^Tt)^2 p\, d\mu}{\alpha} \\ &= \left\langle \chi_{\mathbb{C}}(x^Tt)^2, q\right\rangle. \end{aligned}$$

Hence, we conclude $\nabla g(p) = \chi_{\mathbb{C}}(x^Tt)^2$.

Projection: for any $q \in \mathcal{L}^2(\mathbb{R}^2)$, the projection onto $\mathcal{P}_3$ and $\mathcal{P}_5$ can be calculated by Example 3 and Example 1 respectively. Also, the projection onto $\mathcal{P}_1, \mathcal{P}_2$ and $\mathcal{P}_4$ can be calculated by Example 2. Finally, we approximate the projection onto $\mathcal{P}$ by Dykstra's algorithm described in Theorem 3.

5 Conclusion

In this work, we consider online learning in infinite dimensional space. In particular, we propose an online learning algorithm called online functional gradient descent, which extends the well known online gradient descent algorithm of Zinkevich [25]. We then provide theoretical results for the proposed algorithms. Finally, we illustrate how to apply our algorithm into practical problems in machine learning and operations research, including online classifier selection, risk measure minimization and distributionally robust optimization.

References

[1] Philippe Artzner, Freddy Delbaen, Jean-Marc Eber, and David Heath. Coherent measures of risk. *Mathematical finance*, 9(3):203–228, 1999.

[2] Heinz H Bauschke and Patrick L Combettes. *Convex analysis and monotone operator theory in Hilbert spaces*. Springer Science & Business Media, 2011.

[3] Aharon Ben-Tal, Elad Hazan, Tomer Koren, and Shie Mannor. Oracle-based robust optimization via online learning. *Operations Research*, 0(0):null, 0.

[4] Aharon Ben-Tal and Arkadi Nemirovski. Robust convex optimization. *Mathematics of Operations Research*, 23(4):769–805, 1998.

[5] Dimitris Bertsimas and Melvyn Sim. The price of robustness. *Operations research*, 52(1):35–53, 2004.

[6] Giuseppe Calafiore and Fabrizio Dabbene. *Probabilistic and randomized methods for design under uncertainty*. Springer, 2006.

[7] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine learning*, 20(3):273–297, 1995.

[8] Erick Delage and Yinyu Ye. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations research*, 58(3):595–612, 2010.

[9] L EI-Ghaoui and H Lebret. Robust solutions to least-square problems to uncertain data matrices. *Sima Journal on Matrix Analysis and Applications*, 18:1035–1064, 1997.

[10] Joel Goh and Melvyn Sim. Distributionally robust optimization and its tractable approximations. *Operations research*, 58(4-part-1):902–917, 2010.

[11] Alexandr N Golodnikov, Pavel S Knopov, Panos M Pardalos, and Stanislav P Uryasev. *Optimization in the space of distribution functions and applications in the Bayes analysis*. Springer, 2000.

[12] Elad Hazan. The convex optimization approach to regret minimization. *Optimization for machine learning*, pages 287–303, 2011.

[13] Elad Hazan, Amit Agarwal, and Satyen Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2-3):169–192, 2007.

[14] Elad Hazan, Alexander Rakhlin, and Peter L Bartlett. Adaptive online gradient descent. In *Advances in Neural Information Processing Systems*, pages 65–72, 2007.

[15] George S Kimeldorf and Grace Wahba. A correspondence between bayesian estimation on stochastic processes and smoothing by splines. *The Annals of Mathematical Statistics*, pages 495–502, 1970.

[16] Jyrki Kivinen, Alexander J Smola, and Robert C Williamson. Online learning with kernels. *Signal Processing, IEEE Transactions on*, 52(8):2165–2176, 2004.

[17] Andrew J Kurdila and Michael Zabarankin. *Convex functional analysis*. Springer Science & Business Media, 2006.

[18] Llew Mason, Jonathan Baxter, Peter Bartlett, and Marcus Frean. Boosting algorithms as gradient descent in function space. NIPS, 1999.

[19] Simon Perkins, Kevin Lacker, and James Theiler. Grafting: Fast, incremental feature selection by gradient descent in function space. *The Journal of Machine Learning Research*, 3:1333–1356, 2003.

[20] Andrzej Ruszczynski and Alexander Shapiro. Optimization of convex risk functions. *Mathematics of operations research*, 31(3):433–452, 2006.

[21] Bernhard Scholkopf and Alexander J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, Cambridge, MA, USA, 2001.

[22] Shai Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2011.

[23] Huan Xu, Constantine Caramanis, and Shie Mannor. A distributional interpretation of robust optimization. *Mathematics of Operations Research*, 37(1):95–110, 2012.

[24] Huan Xu and Shie Mannor. Robustness and generalization. *Machine learning*, 86(3):391–423, 2012.

[25] Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the Twentieth International Conference on Machine Learning*, 2003.